WHAT IS CLAIMED IS:

1. A method comprising:
   a) accepting a packet associated with a flow;
   b) generating a flow group identifier from the flow;
   c) determining whether any other packets associated with the flow group are present in a switch fabric;
   d) if it is determined that other packets associated with the flow group are present in the switch fabric, then assigning the packet to a path being used by the flow group, and if it is determined that other packets associated with the flow group are not present in the switch fabric, then assigning the packet to a path using path congestion status information.

2. The method of claim 1 wherein the act of generating a flow group identifier from the flow includes hashing a flow identifier.

3. The method of claim 1 wherein the act of determining whether any other packets associated with the flow group are present in a switch fabric includes maintaining an outstanding packet counter.

4. The method of claim 3 wherein the outstanding packet counter is associated with the flow group identifier.

5. The method of claim 4 wherein the act of maintaining an outstanding packet counter includes incrementing the outstanding packet counter each time a packet belonging to the flow group is sent into the switch fabric, and decrementing the outstanding

packet counter each time a packet belonging to the flow group leaves the switch fabric.

6. The method of claim 5 wherein the act of decrementing the outstanding packet counter is performed in response to a message from an output port.

7. The method of claim 6 further comprising:
- passing the message from the output port to a corresponding input element,
- passing the message from the corresponding input element, through the switch fabric, to another output element, and
- passing the message from the other output element to another input element corresponding to the other output element, wherein the other input element originated the packet.

8. The method of claim 3 wherein the act of maintaining the outstanding packet counter includes resetting the outstanding packet counter if it remains non-zero for more than a predetermined period of time.

9. The method of claim 1 wherein the act of assigning the packet to a path using path congestion status information includes
- selecting a switch plane having at least one uncongested path, and
- selecting an uncongested path of the selected switch plane.

10. The method of claim 9 wherein the act of selecting a switch plane having at least one uncongested path uses a round robin discipline.

11. The method of claim 9 wherein the act of selecting an uncongested path of the selected switch plane uses a round robin discipline.

12. A machine-readable medium having stored thereon a data structure comprising a plurality of entries, each entry including

a) a flow group identifier,

b) an outstanding packet in switch fabric indicator, and

c) a path identifier.

13. The machine-readable medium of claim 12 further including a second data structure comprising a plurality of entries, each entry including

a) the path identifier, and

b) path status information.

14. The machine-readable medium of claim 13 wherein the path status information includes

i) an indicator of whether or not the path has failed, and

ii) an indicator of whether or not the path is congested.

15. Apparatus comprising:

a) an input for accepting a packet associated with a flow;

b) means for generating a flow group identifier from the flow;

c) means for determining whether any other packets associated with the flow group are present in a switch fabric;

d) means for assigning the packet to a path being used by the flow group if it is determined that other packets associated with the flow group are present in the switch fabric, and for assigning the packet to a path using path congestion status information if it is determined that other packets associated with the flow group are not present in the switch fabric.

16. The apparatus of claim 15 wherein the means for generating a flow group identifier from the flow hash a flow identifier.

17. The apparatus of claim 15 wherein the means for determining whether any other packets associated with the flow group are present in a switch fabric maintain an outstanding packet counter.

18. The apparatus of claim 17 wherein the outstanding packet counter is associated with the flow group identifier.

19. The apparatus of claim 18 wherein the means for maintaining an outstanding packet counter increment the outstanding packet counter each time a packet belonging to the flow group is sent into the switch fabric, and decrement the outstanding packet counter each time a packet belonging to the flow group leaves the switch fabric.

20. The apparatus of claim 19 wherein the decrementing of the outstanding packet counter is performed in response to a message from an output port.

21. The apparatus of claim 20 further comprising:

- means for passing the message from the output port to a corresponding input element,

- means for passing the message from the corresponding input element, through the switch fabric, to another output element, and
- means for passing the message from the other output element to another input element corresponding to the other output element, wherein the other input element originated the packet.

22. The apparatus of claim 17 wherein the means for maintaining the outstanding packet counter reset the outstanding packet counter if it remains non-zero for more than a predetermined period of time.

23. The apparatus of claim 15 wherein the means for assigning the packet to a path using path congestion status information include means for
- selecting a switch plane having at least one uncongested path, and
- selecting an uncongested path of the selected switch plane.

24. The apparatus of claim 23 wherein the means for selecting a switch plane having at least one uncongested path use a round robin discipline.

25. The apparatus of claim 24 wherein the means for selecting an uncongested path of the selected switch plane use a round robin discipline.

26. A method for alleviating head-of-line blocking in an input-buffered switch, wherein the switch includes a plurality of

input modules, each input module including virtual output queues and virtual path queues, the method comprising:

a) assigning an incoming cell to an appropriate one of the virtual output queues using cell destination information;

b) providing a head-of-line cell of the one of the virtual output queues to an appropriate one of the virtual path queues using path identifier information of the cell;

c) for an input module-to-switch plane link, selecting one of a number of virtual path queues associated with the link and having at least one cell; and

d) sending the cell from the selected one of the number of virtual path queues over the link.

27. The method of claim 26 wherein the path identifier information of the cell was provided using a dynamic hashing scheme.

28. The method of claim 26 further comprising:

e) determining whether or not the cell sent over the link was the last cell of a packet; and

f) if it was determined that the cell sent over the link was the last cell of a packet, then instructing the virtual output queue to send cells of a next packet to an appropriate one of the virtual path queues.

29. For use in a switch, an input module comprising:

a) a plurality of virtual output queues for accepting cells; and

b) a plurality of virtual path queues for accepting head-of-line cells from the plurality of virtual output queues.

30. The input module of claim 29 wherein the number of the plurality of virtual output queues equals a number of output ports of the switch.

31. The input module of claim 29 wherein the number of the plurality of virtual path queues equals a number of paths through a switch fabric of the switch.

32. The input module of claim 29 wherein the number of the plurality of virtual path queues equals a product of (a) a number of switch planes of a switch fabric of the switch and (b) a number of paths through each of the switch planes.